A Conversational Agent for the Improvement of Human Reasoning Skills

Laura Wartschinski, Nguyen-Thinh Le, Niels Pinkwart

Abstract: Human reasoning is the ability to make sound and goal-oriented decisions and is therefore highly relevant in daily life. However, its importance has not yet been addressed explicitly in education. In this paper, we propose to develop a computer-based conversational agent named Liza for improving human reasoning skills. Liza is able to hold conversations with humans to help them solve a small selection of well-studied reasoning problems. Such a conversational agent allows a much more scalable and inexpensive approach for teaching at least basic reasoning principles. The evaluation study shows that the conversational agent Liza improved the reasoning skills of the participants, who had conversations with the agent to solve reasoning problems and that the group using Liza achieved much better learning effects than a group studying with a non-interactive online course that was implemented as a control condition.

Keywords: human reasoning education, pedagogical agent, dialogue system, natural language processing, heuristics and biases

1   Introduction

People are behaving in an instrumentally rational if they pick whatever behavior is best suited to reach their goals, no matter what those goals may be [St12, p.345]. In order to choose the right behavior, it is often necessary to form correct beliefs about the world. Psychological research of the last three decades has shown that people's decisions often deviate from the normatively correct behavior, which is interpreted as irrationality in humans [SW00]. There are many examples for people who behave irrationally, e.g. in displaying errors in calibrating their degrees of belief, not processing information correctly, or failing at assessing probabilities [EO96]. This is mainly because people tend to rely on heuristics that lead to systematic errors, also called cognitive biases [TK74, p. 1124], even when they possess the necessary skills to solve a task correctly [Ka03; Kl01]. The Dual Process Theory offers an explanation for this behavior [ES13; Ka03]. Furthermore, it is important to note that rationality seems to be a compound quality, as performance on different reasoning tasks shows significant correlation [SW00, p.664], and is also distinct from general intelligence and cognitive ability [SW00; St09]. It has been found that teaching can improve reasoning skills [La04], but in order to do so, requires both repeated explanations and the use of many examples [Ch86, p.293]. Here, the use of educational software may offer benefits. It has been shown that software can successfully be used to train general cognitive skills [Ha11], and educational games [Mo15] as well as videos and learning software [Ge00, p.132] have already been used to improve their user's reasoning. Pedagogical agents, however, who have proven

themselves to be useful as learning motivators and guides [KB07], have not been applied on this subject so far. However, the applicability of many reasoning skills to everyday life predestines the topic for learning in a dialogue-like environment as it is provided by a pedagogical agent. Also, there are not many human teachers available for the field. Therefore, the application of pedagogical agents as teachers for reasoning seems worthwhile. In this work, we prove that a pedagogical agent can improve its users' performance on classical reasoning tasks. Furthermore, we show it has a stronger impact on their performance than a non-interactive online course on the same subject.

2   State of the Art

The goal of this work is to create an Intelligent Tutoring System (ITS) for the teaching of reasoning. ITS are computer programs that interact with students, guide them through the process of learning and provide personalized instructions, tasks and feedback to them. They strive to improve the learning of their students by delivering adaptive and personalized content [VR14]. ITS can be accompanied by conversational agents that engage in chat-oriented or task-oriented dialogues with the user [BL12]. The software system developed in this work also aims to fulfill the role of a Pedagogical Agent, that is, a computer-based agent for educational purposes that delivers information, teaches its users, reacts to its environment (specifically the user input) and chooses the appropriate actions. Current examples for already realized pedagogical agents with similar goals can be found in [Sh09], [ER14] and [La13]. None of those agents however was so far aimed at teaching a general cognitive skills like reasoning. In fact, researchers have only started to examine how modern media could be used in teaching reasoning. For teaching critical thinking, a related composite skill covering argumentation, informal logic and correct judgment, it has been found that computer-assisted training has been superior to training that was carried out traditionally [Hi03, p.188; Ge01, p.547]. Psychologists have already expressed an explicit wish for an interactive software as a teaching tool for this subject [Wa00, p.37].

Previous work on pedagogical agents as tutors for other areas shows very promising result. The agents engage students via social interactions and encourage them to invest effort and persist in learning [KB07], motivating them and keeping them interested [Mo01]. They are able to increase the pleasantness of the learning experience so that user enjoy working with them [Le97]. In general, the use of a pedagogical agent as a social model to improve the motivation and attitude of the student towards the subject is especially effective and makes students face challenges, put effort in their work and persist in learning [KB07]. Students that are taught by agents spend more time with their tasks and easily acknowledge their own mistakes [Ch09]. Those benefits are expected to also work in favor of a pedagogical agent for teaching reasoning and are now applied to this subject for the first time.

3 Design, Content, and Implementation

3.1 Design

This works aims to create a conversational agent, called Liza, which provides natural and, to a certain degree, free communication between computer and user. A text-only approach was chosen to invoke a stronger social presence [DMA05, p.13] and avoid irritations and too high expectations caused by animations that might disappoint the user [Gu11]. Chat-like text pieces and additional images for illustration purposes are therefore the chosen means of communication.

The agent provides explanations on seven topics in the field of reasoning, heuristics and biases, and guides the user through a series of questions, reviewing their level of understanding and correcting their mistakes while at the same time maintained a somewhat natural conversation. It is task-oriented, but uses human social behavior and emotions to improve the learning experience and results. Since effective pedagogical agents suggest correct solutions, provide hints and explanations, give examples, reference to relevant background material and also test the student's abilities, the agent described here implements all those requirements. It aims for a mentor-like role [BK05; Ba00] as a trustworthy and guiding expert that is nevertheless on the same level as the student. Furthermore, the proposed agent uses small talk about personal preferences and sometimes jokes to lighten the mood, as this has been shown to reduce stress [MAM98] and increase enjoyment [BSY09] and a good relationship between agent and user [Gu11; Ku07]. The agent refuses to engage in conflict and stays professional and calm. It does display ‚emotions', however. Emotionally responsive agents improve self-efficacy and interest in students [KB07], and the use of emotion makes the interaction much more natural, even when only basic emotions are used [Pi03]. Therefore, Liza shows various sentiments, including pride, joy, curiosity, interest, playfulness, disappointment, sadness and even insecurity, with an emphasis on positive emotions. The agent displays a clear interest in the student's learning success and frequently praises and encourages the users [RN96, p. 310; Sh09]. Praise and flattery expressed by the agent can have a very powerful effect, making the user feel motivated, capable and happy [FN97]. The agent's persona is that of an inexperienced piece of software that tries to learn how to be a good teacher. This narrative allows for a more laid-back style of conversation and gives the agent a plausible goal and also a convenient excuse for any errors.

The agent starts out with a greeting phase in which it explains the process and starts to learn things about the user. Then, during the main part of the program, fourteen questions from the seven different reasoning tasks are presented (two from each). When a task category appears for the first time, an introduction is given, which is accompanied by helpful images and options to give more in-depth explanation. Then the question is asked. For each such question, Liza gives some connecting transition that leads naturally to the topic, then describes a situation and finally asks the question. Knowledge about the user's preference is used to personalize the stories and introductions incorporating the

student's hobbies or personality traits. The user's answer can be decisions, picking several from different options, giving reasons, statements or numbers. If the answer was correct, the agent displays joy and is proud of the user, if not, it might question its own abilities as a teacher and offer to explain the problem again. The user's confidence in their own answers is also determined and later used to provide feedback about their calibration. In between the questions, there are several options for off-task chat. In the end, the agent provides a detailed evaluation of the user's success.

An examplary piece of dialog out of the middle of die dialogue can look like the following. Liza: *"Are you familiar with the sunk cost fallacy?"* User: "*Never heard of that."* Liza: *"Wonderful! This is my chance for being a good teacher!"* (Liza goes on an explains the concept, then makes sure that the user understood it, then asks the question. The user gives a wrong answer) Liza: *"Hm, I think you might be wrong this time. Do you want me to tell you why?"* User: *"Yeah, go ahead"* (Liza explains the solution, the user agrees) Liza: *"By the way, was that a hard question for you?"* User: *"Probably. I have to wrap my head around it."* Liza: *"Thank you, I noted that down. It's a pleasure to work with you!"* User*: "I like you."* Liza: *"That's so nice! Oh, the next question is about your favority city, Amsterdam. Are you ready?"*

The conversation between user and agent is created as a mixed-initiative dialogue. The agent does not wait forever for a user to answer, but uses flexible waiting times similar to the way a human would act in a conversation. It starts asking worried questions if the user does not react for a very long time, and it offers to give a hint if the answer takes a while. It also reacts to inputs at any time, even if the user just chimed in an insult or a question while the agent was talking. However, Liza does not offer completely free off-task conversation, but rather just answers with short but appropriate messages to whatever was said. Still, those reactions seem sufficiently natural to correspond to the image of a teacher that is largely concentrated on the task. During the whole process, the user can always intervene, for example to stop the dialogue entirely. The direction the conversation takes is changed dynamically to accord for the user's progress and results. The agent picks its replies and questions from a sufficiently large database that stores ten to twenty different variants of reactions for every situation to ensure a natural feeling for the conversation without repetition.

3.2   Content

The content of the agent was taken from the most well-known reasoning tasks for which studies already show that improvement through teaching and exercises is feasible. The seven topics are: Bayesian reasoning [Be81; HG98], the Law of Large Numbers [TK74], the Gambler's Fallacy [Kl01], Wason's Selection Task [Wa68; Kl01], Covariance Detection [St09; Kl01], the Sunk Cost Fallacy [LMN90; Th80] and Belief Bias in Syllogistic Reasoning [EBP83; MN89]. To give one example: The sunk cost fallacy is the irrational desire to stick with the outcomes of previous decision because something was already invested (the ,sunk cost'), even if the continuing commitment provides less

benefit than another option. One example task used by the agent would be started with a short intro about a trip the user takes to a city he likes (inserting the user's favorite city, if known to the agent), and then explaining how he went to a local cinema and paid a lot for the ticket. The agent states that the movie is very boring, that the user is also the only person watching it. Liza describes that the user could spend their time better by doing some other activity. Now it's the user's turn to decide. If the answer is to stay in the cinema because the ticket was costly and is already paid, this is counted as incorrect and an example of the sunk cost fallacy. If the user decides to leave the cinema, this is counted as correct.

3.3   Implementation

Liza was programmed from scratch in Java and was developed using an object-oriented approach. It consists of several units. The Control Unit is responsible for the general management of the dialogue—it keeps track of the state of the conversation, decides what to do next and interacts with all the other components. Input and output are handled by the UI in form of a typical chat interface. The input by the user is then passed on to the Parser, which determines the content or meaning and returns that to the Control Unit, which in turn can decide to react with a phrase from the Phrase Base. The stories are retrieved from the Story Store. A story is a task that is to be presented to the user, complete with descriptions, questions, reactions, explanation and everything else that is necessary for processing. While the Control Unit is executing the stories, correct and incorrect responses are stored by the User Evaluation.
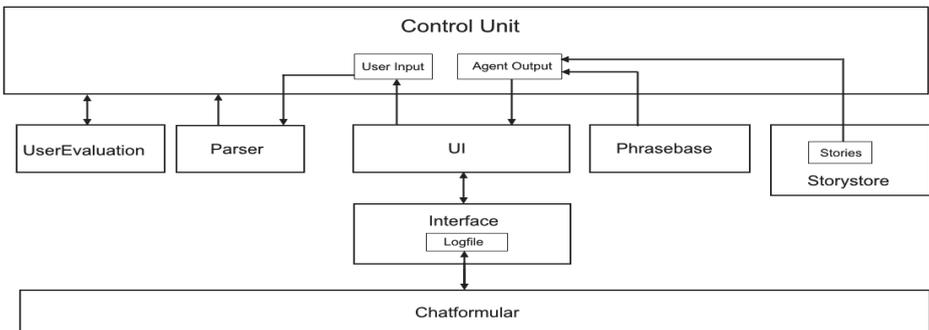

Fig. 1: simplified software architecture

The parsing module relies on basic pattern matching, but is highly specialized for its core purpose: to assess whether the specific questions asked by the agent were answered correctly. Every input is parsed with regard to the context of the message, and more than 80 different context types are distinguished. Knowing the specific context, the agent can be relatively sure that a certain keyword translates to a right or wrong answer to the question asked, whereas without the context, the agent would be clueless because of its rather simple parsing system. This approach reduces the complexity of the parsing a lot

and can be seen as a shortcut to achieve more human-like results with simple methods. The input is parsed recursively with respect to phrases that alter the meaning, like "No" or "wouldn't", looking for specific keywords. The combination of keywords that are found and that point to a correct, incorrect or undecided answer determines the final interpretation of the input. Of course, the agent will ask for clarification if the parsing failed. If the user switches the context to something not foreseen by the agent's programming, it will not be able to understand it. For example, if a user asked the agent about the weather, it would not 'understand' it, but, depending on the context, glance over it or ask for a rephrasing fitting to the question asked. Partly, this problem is softened by general answers that are interpreted as fitting to the context even if they are not ("don't you think I should be asking the questions?"). The described parsing mechanism is used for any user input. Since the agents always know the context, it can look out for numbers or verbal statements of quantities or probabilities ("almost certain!") after asking for the likeliness of a result, and look for affirmation or rejection after asking a yes/no question. That way, the parser can have a very good guess on what the user wanted to communicate.

4 Evaluation

4.1 Study Design

To evaluate the impact of the agent, the following hypotheses were tested: 1. The participants do better on the reasoning tasks after they practiced with the agent than they did before, suggesting they improved their reasoning skills. 2. Talking to the agents leads to a stronger improvement than the non-interactive online course, therefore in the second test, the performance of the participants who talked to the agent will be significantly better than the performance of the participants who took the online course.

The survey for evaluating the success of the agent was conducted using a privately set up server and php and html forms. For each participant, some general information was collected: age, gender, degree of education and level of English skills. The following study design was used: The test subjects were randomly distributed in a treatment group and a control group. The treatment group got to talk to the agent using an online chat interface. The control group read a short text about each of the biases tasks and learned about the underlying principle of reasoning. The texts were taken from different text books and online courses used in actual teaching. For each participant, a test before the intervention and another one afterwards were administered and the performance for every single task was collected. The subjects were forced to pick an answer even if they were not totally sure. This was done to assess even slight preferences to one pick, despite subjects being not totally sure about their reply. For every topic that was assessed with a task, the mean performance before and after the treatment was calculated. For every task, a correct solution was evaluated to have the value 1, an incorrect answer was assigned the value 0. The instructions read as following: "Please solve the following questions by

giving the answers that seem most reasonable, sensible or logical. It's perfectly normal that some of the answer may seem obvious (they probably are—there's no mean trick behind them), and others are harder. Feel free to use a calculator or make notes if it helps you, but please solve the questions on your own. In general, just trust your good reason!" The questions themselves were either the exact same or, if this was not possible, closely modeled after the questions used in current research on reasoning. Participants were obtained by inviting acquaintances and fellow students. 25 more subjects however were paid participants from an online platform called Clickworkers, and 14 were paid participants from various online communities that were paid with $3 gift cards. In total 65 test subjects completed the survey. The average age of participants was 28.9 years, 44.6% were female, and 82% were fluent or native English speakers. Roughly half of the participants had a Bachelor's degree or a higher education, which is in part due to the acquisition process.

4.2   Results

A two-sample test for proportions was conducted for the treatment group (that was talking to the agent), to test the hypothesis that performance after the treatment would be different to the performance before the treatment. The same test was done for the control group. Also, the change in performance was determined, comparing the mean before the test to the mean afterwards, still assigning a "1" to a correct and a "0" to an incorrect result. For the treatment group, five of the seven tasks showed a significant or very significant improvement, while the control group only improved significantly in one of the tasks. The five significantly ($p < 0.05$) or highly significantly ($p < 0.01$) improved tasks and their improvements were: Sunk Cost Fallacy, Gambler's Fallacy, Bayesian Reasoning, Regression to the Mean and Covariance Detection (see Table 1).

Furthermore, a t-test was conducted comparing the results of the treatment group in the second test (after talking to the agent) to the control group in the second text (after reading their texts), to see if there was any significant difference between them. As it turned out, the treatment group was significantly or very significantly better in the second test than the control group in four of the tasks. There was no task in which the control group performed better in the second task than the treatment group, regardless of significance. A linear regression analysis was performed and showed than none of the demographic variables had any influence on the improvements gained through working with the agent or taking part in the control group's online course. Not even the level of English accounted for any significant change in the effectiveness of the interventions.

Feedback from participants suggested that they perceived the study as interesting. Some gave feedback like "The chatbot is really nice :)" or "Thank you for this interesting experience, I really liked Liza and her funny and helpful way of teaching. :)". Those answers suggest that the people giving them had a positive social experience with the conversational agent. Also, some participants expressed ideas on how to further improve the agent. During the interaction, the users did mainly focus on the dialogue structure the

agent provided and showed no sign of inappropriately high expectations. They stuck to the tasks they were given by Liza and followed her guidance easily.

Task Category	Cond.	Treatment		Control	
		Mean	Std. Err.	Mean	Std. Err.
Sunk Cost Fallacy	before	0.80	0.07	0.63	0.09
	after	1.00	0.00	0.73	0.08
	differ.	**0.20****	0.07	**0.10**	0.12
Gambler's Fallacy	before	0.83	0.06	0.70	0.08
	after	0.97	0.03	0.83	0.07
	differ.	**0.14***	0.07	**0.13**	0.11
Bayesian Reasoning	before	0.17	0.06	0.07	0.05
	after	0.31	0.08	0.23	0.08
	differ.	**0.14***	0.07	**0.17***	0.09
Belief Bias in Syll. Reasoning	before	0.74	0.07	0.57	0.09
	after	0.63	0.08	0.53	0.09
	differ.	**-0.11**	0.11	**-0.03**	0.13
Regression to the Mean	before	0.34	0.08	0.50	0.09
	after	0.77	0.07	0.63	0.09
	differ.	**0.43****	0.11	**0.13**	0.13
Covariation Detection	before	0.71	0.08	0.50	0.09
	after	0.89	0.09	0.53	0.09
	differ.	**0.17***	0.11	**0.03**	0.13
Wason's Selection Task	before	0.40	0.08	0.37	0.09
	after	0.57	0.08	0.30	0.08
	differ.	**0.17**	0.12	**-0.07**	0.12

Tab. 1: comparison of treatment (talking to the agent) vs. control (online course) group

4.3   Discussion

This agent was created hoping that the interaction with it would improve the ability of the participants to solve the reasoning tasks. For five tasks, there was a significant increase in performance. The two remaining tasks had no significant effect, although the performance on the selection task did noticeably improve. Therefore, the first hypothesis is largely confirmed. The work with the agent also proved itself to be more efficient than the control group in many cases. The gains in performance were stronger and significant in five cases, whereas the control group could only significantly improve in one task. It can therefore be assumed that the agent provided a more efficient intervention, confirming the second hypothesis.

Only in bayesian reasoning, the control group improved more than the treatment group, but both effect sizes were very close to each other and the task was generally difficult to

teach. Both tasks with the more disappointing, non-significant results (Bayesian Reasoning and Selection Task) are already considered belonging to the hardest reasoning tasks, having generally very low rates in being solved correctly, and being hard to explain, it can be assumed that a short training was not enough to convey the intended message. After all the subjects spent on average 35-45 minutes with the agent, but this time had to be split between the seven different topics they learned about and the general explanation and small talk parts, leaving only roughly 5 minutes on average per topic.

More participants than the acquired 65 would of course offer a more accurate picture of the observed effects. A restriction for the execution of the study was the use of the English language by the bot, which required the participants to be able to write and read in English. The study was conducted online, allowing for participants from different locations to take part in it. It seems highly likely that this allowed for more participants to be acquired, but the sample was not representative. A bias towards younger people is inevitable. Furthermore, many of the participants were motivated extrinsically via payment, in contrast to a real-world application of the agent in a teaching setting without monetary reward. Probably the strongest limitation was the harsh time limit, as none of the participants could be expected to spend more than an hour on the whole process. Therefore, the interaction with the agent was limited in duration, restricting the time for teaching, and it was not possible to assess every task in the test before and after the intervention with a whole set of questions (instead of only one), since that would have further increased the duration of the study. It would have been very interesting to analyze the long-term effects, but only six participants signed up for a follow-up test some weeks later. They all performed very well, but this small sample is not significant in any way.

5   Conclusion and Future Work

The contribution of this paper is two-fold. First, we presented a conversational agent that was designed to improve skills for a selection of core reasoning tasks and to the best of our knowledge, this learning domain has not been addressed in the community of technology-enhanced learning. The second contribution is the empirical study of the developed conversational agent for human reasoning. Although positive impact could not be found in solving all the tasks, maybe because Wason's Selection Task and Bayesian reasoning are both very complicated concepts to teach, the evaluation study showed very satisfying results with most of the tasks. The agent demonstrated to be clearly more effective than participating an online course. It can be concluded that performance on solving certain reasoning tasks can indeed be improved with the help of pedagogical agents, which is especially fortunate as there are not many teachers in this area, despite its importance for arguably everybody's personal life.

Given the nature of the agent, it can always be improved by adding more special parsing options and appropriate responses. By extending the scope of correctly parsed input, the

agent can become even more convincing, although there might be diminishing returns for increasingly rare topics. More elaborate dialogue patterns could be added, too. Further improvements could be made by allowing the user to ask the agent for more questions of a certain kind, or the agent suggesting more training in an area where the user seems to lack understanding. The agent could probably be put to a more effective use if it was used not only for a single intervention, but regularly. Storing the past performances of users, remembering their preferences, and specifically targeting the areas where they can improve come to mind as obvious possibilities for expanding the agent. Also, right now the agent does only teach students in seven limited and strictly restricted topics. Improvements could include more topics (e.g. addressing some more of the wide variety of known fallacies and biases), offering a more general approach (e.g. giving some philosophical background on knowledge, its representation and a more refined approach on what rationality means), or digging into new areas of the field (e.g. the area of 'critical thinking' which includes open-mindedness and discussing the quality of arguments and which was originally intended to be part of the agent, but had to be left out due to time and scale concerns). By those means, the 'curriculum' of the agent could be extended in many directions. First and foremost, it would be very interesting to validate the results with a much larger sample than the 65 participants who were available. An expansion of this research could be a study over the course of several weeks or months, to have the users spend much more time with the agent, to see if this approach is more effective. Long term effects could be investigated by reassessing user's performance after weeks, months or even years.

References

[Ba00]   Baylor, A.: Beyond butlers: Intelligent agents as mentors. Journal of Educational Computing Research 22/4, S. 373–382, 2000.

[Be81]   Berwick, D.M. et al.: When doctors meet numbers". The American journal of medicine 71/6, S. 991–998, 1981.

[BK05]   Baylor, A.; Kim, Y.: Simulating instructional roles through pedagogical agents. Int. Journal of Artificial Intelligence in Education 15/2, S. 95–115, 2005.

[BL12]   Banchs, R.E.; Li, H.: IRIS: a chat-oriented dialogue system based on the vector space model. In: Proceedings of the ACL System Demonstrations, 2012.

[BSY09] Bickmore, T.; Schulman, D.; Yin, L: Engagement vs. deceit: Virtual humans with human autobiographies". In: Int. Workshop on Intelligent Virtual Agents. S. 6–19, 2009.

[Ch09]   Chase, C.C. et al: Teachable agents and the protégé effect: Increasing the effort towards learning. Journal of Science Education and Technology 18/4, 2009.

[Ch86]   Cheng, P.W. et al.: Pragmatic versus Syntactic Approaches to Training Deductive Reasoning. Cognitive Psychology 18, S. 293–328, 1986

[DMA05] Dirkin, K.H.; Mishra, P.; Altermatt, E.: All or nothing: Levels of sociability of a

pedagogical software agent and its impact on student perceptions and learning. Journal of E. Multimedia and Hypermedia 14/2, 2005.

[EBP83] Evans, J.; Barston, J.L.; Pollard, P.: On the conflict between logic and belief in syllogistic reasoning. Memory & Cognition 11/3, S. 295–306, 1983.

[EO96] Evans, J.; Over, D.E.: Rationality and Reasoning. Psychology Press, 1996.

[ER14] Epstein, D.; Reategui, E.: A pedagogical agent with embedded data mining functions to support collaborative writing. Learning 12/13, S. 14, 2014.

[ES13] Evans, J.; Stanovich, K.E.: Dual-process theories of higher cognition advancing the debate". Perspectives on psychological science 8/3, 2013.

[FN97] Fogg, B.J.; Nass, C.: Silicon sycophants: the effects of computers that flatter. Int. Journal of Human-Computer Studies 46/5, S. 551–561, 1997.

[Ge00] Gelder, T. van: Learning to reason: A Reason!Able approach". In: Proc. of the Fifth Australasian Cognitive Science Society Conference. Adelaide: Causal, 2000.

[Ge01] Gelder, T. van: How to improve critical thinking using educational technology. In: Proc. of the 18th Annual Conference of the Australasian Society for Computers in Learning in Tertiary Education, S. 539–548, 2001.

[Gu11] Gulz, A. et al.: Building a Social Conversational Pedagogical Agent. In: Conversational Agents and Natural Language Interaction: Techniques and Effective Practices, S. 128, 2011.

[Ha11] Haferkamp, N. et al.: Training disaster communication by means of serious games in virtual environments. Entertainment Computing 2/2, S. 81–88, 2011.

[HG98] Hoffrage, U.; Gigerenzer, G.: Using natural frequencies to improve diagnostic inferences. Academic medicine 73/5, S. 538–540, 1998.

[Hi03] Hitchcock, D.: The effectiveness of computer-assisted instruction in critical thinking. Informal Logic: Reasoning and Argumentation in Theory and Practice 24/3, S. 183–217, 2003.

[Ka03] Kahneman, D.: A perspective on judgment and choice: mapping bounded rationality. American psychologist 58/9, 2003.

[KB07] Kim, Y.; Baylor, A.: Pedagogical agents as social models to influence learner attitudes. Educational Technology 47/1, S. 23–28, 2007.

[Kl01] Klaczynski, P.A.: Analytic and Heuristic Processing Influences on Adolescent Reasoning and Decision-Making. Child development 72/3, S. 844–861, 2001

[Ku07] Kumar, R. et al.: Supporting students working together on math with social dialogue. In: In Workshop on Speech and Language Technology in Education, S. 96–99, 2007.

[La13] Lane, H.C. et al.: The effects of a pedagogical agent for informal science education on learner behaviors and self-efficacy. In: Int. Conference on Artificial Intelligence in Education, S. 309–318, 2013.

[La04] Larrick, R.P.: Debiasing. In: Blackwell handbook of judgment and decision making, S. 316–337, 2004.

[Le97]   Lester, J.C. et al.: The persona effect: affective impact of animated pedagogical agents. In: Proc. of the ACM SIGCHI Conference on Human factors in computing systems, S. 359–366, 1997.

[LMN90] Larrick, R.P.; Morgan, J.N.; Nisbett, R.E.: Teaching the use of cost-benefit reasoning in everyday life. Psychological Science 1/6, S. 362–370, 1990.

[MAM98] Mulken, S.van; André, E.; Müller, J.: The persona effect: how substantial is it?" In: People and computers XIII. Springer, S. 53–66, 1998.

[MN89]   Markovits, H.; Nantel, G.: The belief-bias effect in the production and evaluation of logical conclusions. Memory & Cognition 17/1, S. 11–17, 1989.

[Mo01]   Moreno, R. et al.: The case for social agency in computer-based teaching: Do students learn more deeply when they interact with animated pedagogical agents? Cognition and instruction 19/2, S. 177–213, 2001.

[Mo15]   Morewedge, C.K. et al.: Debiasing Decisions Improved Decision Making With a Single Training Intervention. In: Policy Insights from the Behavioral and Brain Sciences 2/1, S. 129-140, 2015.

[Pi03]   Picard, R.W.: Affective computing: challenges. Int. Journal of Human-Computer Studies 59/1, S. 55–64, 2003.

[RN96]   Reeves, B.; Nass, C.: How people treat computers, television, and new media like real people and places. CSLI Pub. and Cambridge University Press, 1996.

[Sh09]   Shen, E.: Effects of agent emotional support and cognitive motivational messages on math anxiety, learning, and motivation, 2009.

[St09]   Stanovich, K.E.: What intelligence tests miss: The psychology of rational thought. Yale University Press, 2009.

[St12]   Stanovich, K.E.: On the distinction between rationality and intelligence: Implications for understanding individual differences in reasoning. In: The Oxford handbook of thinking and reasoning, S. 343–365, 2012.

[SW00]   Stanovich, K.E.; West, R.F.: Individual differences in reasoning: Implications for the rationality debate? Behavioral and brain sciences 23/5, 2000.

[Th80]   Thaler, R.: Toward a positive theory of consumer choice. Journal of Economic Behavior & Organization 1/1, S. 39–60, 1980.

[TK74]   Tversky, A.; Kahneman, D.: Judgment under uncertainty: Heuristics and biases. Science 185/4157, S. 1124–1131, 1974.

[VR14]   Veletsianos, G.; Russell, G.S.: Pedagogical agents. In: Handbook of research on educational communications and technology. Springer, S. 759–769, 2014.

[Wa00]   Walton, D.: Problems and useful techniques: My experiences in teaching courses in argumentation, informal logic and critical thinking. Informal Logic 20/2, 2000.

[Wa68]   Wason, P.C.: Reasoning about a rule. The Quarterly journal of experimental psychology 20/3, S. 273–281, 1968.